

A Markov chain model of evolution in asexually reproducing populations: insight and analytical tractability in the evolutionary process

Daniel Nichol¹, Peter Jeavons¹, Robert Bonomo², Philip K. Maini³, Jerome L. Paul⁴, Robert A. Gatenby⁵, Alexander R.A. Anderson⁵ & Jacob G. Scott^{3,5}

¹ Department of Computer Science, University of Oxford, Oxford, UK

² Department of Medicine, Louis Stokes Department of Veterans Affairs Hospital, Cleveland, OH, USA

³ Centre for Mathematical Biology, University of Oxford, Oxford, UK

⁴ School of Computing Sciences and Informatics, University of Cincinnati, Cincinnati, OH, USA

⁵ Integrated Mathematical Oncology, H. Lee Moffitt Cancer Center and Research Institute, Tampa, FL, USA

Correspondence:

Jacob G Scott, e-mail: jacob.g.scott@gmail.com and

Daniel Nichol, e-mail: daniel.nichol@st-annes.ox.ac.uk

Keywords: theoretical biology, drug resistance, bacteria, mathematical model, evolution, cancer

Abstract

The evolutionary process has been modelled in many ways using both stochastic and deterministic models. We develop an algebraic model of evolution in a population of asexually reproducing organisms in which we represent a stochastic walk in phenotype space, constrained to the edges of an underlying graph representing the genotype, with a time-homogeneous Markov Chain. We show its equivalence to a more standard, explicit stochastic model and show the algebraic model's superiority in computational efficiency. Because of this increase in efficiency, we offer the ability to simulate the evolution of much larger populations in more realistic genotype spaces. Further, we show how the algebraic properties of the Markov Chain model can give insight into the evolutionary process and allow for analysis using familiar linear algebraic methods.

Introduction

Understanding the evolutionary trajectories of populations under given selective pressures is a fundamental problem in biology. In particular, understanding which traits are likely to be selected for within a population can help us better treat disease and understand ecological changes. Traditional computational models of evolution have often simulated populations explicitly (Moran et al. [1962]) - keeping a population of members and probabilistically iterating phases of reproduction, mutation and selection. These methods are cumbersome and can become too computationally complex to simulate for large populations or large numbers of accessible genotypes.

Past successful simplifications of this evolutionary process for asexually reproducing organisms have been to model a genotypically homogeneous population which undertakes a stochastic walk of mutations, mutating at each time step to a fitter variant as a population. This model makes the so-called Strong Selection Weak Mutation assumption and has been used to study which evolutionary trajectories are inaccessible to a population of organisms and how these trajectories are changed by sign epistasis - the situation in which a given mutation may be beneficial under certain selective pressures but deleterious under others (Weinreich et al. [2005, 2006], Tan et al. [2011], Poelwijk et al. [2011]). Certainly this model is efficient to simulate, but it tells us little about population dynamics. In particular, how do different selective pressures influence how a heterogeneous population diversifies or converges over time?

In this paper we present two models of evolution: one an explicit stochastic model which tracks each individual within the population and simulates mutation and selection, and the other a more abstract description using a time-homogeneous Markov chain. We show that the two models, given the same evolutionary landscape and initial population, result in nearly identical populations. Further, once the evolutionary process has been encoded in a Markov chain, we show that there is additional insight which can be gained by familiar linear algebraic analysis. Finally, we show that use of this novel method of encoding the evolutionary process is far less computationally intensive than more standard explicit models and hence allow for exact calculation of evolution for large populations on large, rugged evolutionary landscapes.

Fitness Landscapes and the Genotype-Phenotype Map

The mapping from genotype to phenotype is a familiar concept, but one that has eluded rigorous treatment in the genomic era. While mutation certainly occurs at the level of the genotype, selection operates at the level of the phenotype, making this mapping central to any study of evolution. We will begin the process of modelling this mapping by utilizing the representation of the genotypes of

an asexually reproducing organism as presented by Weinreich and colleagues to study evolutionary trajectories (Weinreich et al. [2005]). Once we have established a framework for the genotype, which constrains the allowable paths through mutation space, we will invoke a genotype-phenotype mapping to establish the ‘forces’ of selection. This two-level system will then constitute the basis for our models.

To this end, we represent the genotype of an organism by a bit string of length N and model mutation as the process of flipping a single bit within this string. This gives a set of possible genotypes of size 2^N in which each genotype, x , has N one-mutation neighbours - precisely those genotypes for which the Hamming distance (Hamming [1950]) from x is 1. As such, our genotype space (\mathcal{G}), can be represented by an undirected N -cube graph with vertices which represent genotypes, and edges which connect neighbours at Hamming distance 1 (See Fig. 1a).

We then define a selective pressure on our graph that drives evolution, for example through an environmental change or drug application, as a fitness function which acts on each vertex in the graph, \mathcal{G} ,

$$f : \{0, 1\}^N \rightarrow \mathbb{R}^{\geq 0}. \quad (1)$$

This fitness function represents a genotype-phenotype map in the simplest sense - assigning to each genotype a single real-valued fitness. This could be, for example, thought of as resistance to an applied drug, as it was by Weinreich (Weinreich et al. [2006]) and later Tan (Tan et al. [2011]), to study evolutionary trajectories of *E. Coli* under selection by different beta-lactam antibiotics.

Using these fitness values we construct a directed evolutionary graph on the set of 2^N possible genotypes where there exists an edge from a to a neighbour b if and only if $f(b) > f(a)$ (See Fig. 1b). This graph provides a model of which mutations and which evolutionary trajectories (series of mutations) are possible - those which increase the fitness of the individual at each step.

An Explicit Computational Model of Evolution

Our first attempt to the evolutionary dynamics of this system was by an explicit simulation of a population undergoing mutation and selection over time, similar to the familiar Moran and Wright-Fischer processes ([Moran et al., 1962, Fisher, 1930]). We keep a population of n individuals and at each time step replace each individual with k of its fitter neighbours. From the nk individuals we have, we choose the fittest n to survive the selection phase. We repeat this process until each individual in the population has no fitter one-mutation neighbours. Figure 2 shows this process for $n = 4$ and $k = 2$. This model, whilst certainly simplified, attempts to mimic the process




Figure 1: (a) The genotype space, \mathcal{G} , of an $N = 3$ cell, where the genotypes are represented by bit strings of length N and the edges of the graph represent Hamming distance 1 connections. (b) An example phenotype graph, \mathcal{P} , with vertices represented by integer fitness values determined by the mapping, f , from equation 1, and directional edges representing of evolutionarily allowable mutational transitions.

of evolution as closely as possible, so we might expect it to make good predictions about the evolutionary trajectories of real-world organisms.

This algorithm requires, at each iteration nk , random samples from a uniform distribution. We know that in a random landscape the average mutational path length to a fitness optima is $\log_2(N-1)$ (Altenberg [1997]). Hence, the expected number of random samples required is $\mathcal{O}(nk \log_2(N-1))$ well as the extra time associated with the sorting in the selection step, which we choose to ignore as it is considerably quicker than the random sampling.

As the genotype spaces of real-world organisms can be very large, this algorithm can require many iterations for the population to stabilize into a final population and to fully explore a large and multi-peaked fitness landscape we would need to maintain a large population. Further, this algorithm must be run many times to give an expected distribution for which the error is acceptable. For example, the simulation required to produce Figure 3 required 1000 iterations on a landscape for which $N = 5$. The running time of this algorithm makes it intractable for finding expected distributions in genotype spaces which are large enough to be useful.

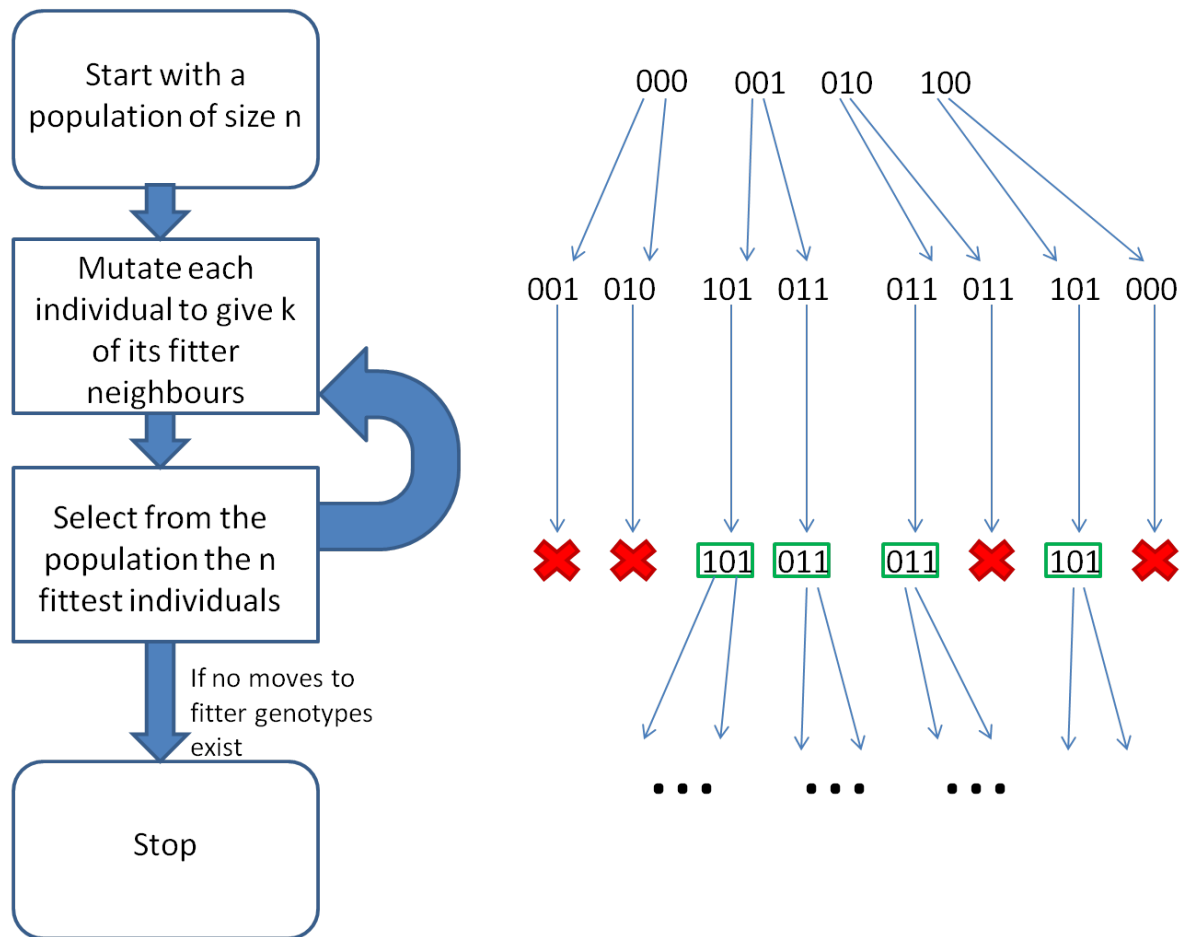


Figure 2: An explicit stochastic evolutionary process. Each individual in a population is mutated at a baseline mutation rate to each of its k fitter neighbours by a stochastic process biased by the difference in fitness. Selection (green rectangle) acts on the next generation by eliminating the least fit individuals (red X). The process repeats itself until there are no fitter neighbours to which to mutate.

A Markov Model of Evolution

To overcome the problem of inefficiency in explicit simulations and, in particular, the need to iterate them many times to obtain significant results, we will build an algebraic model of the same evolutionary process. We represent the genotype of an organism as a bit string and the phenotype as a specific mapping of each vertex, as before, and model evolution of an organism within a population as an uphill stochastic walk in phenotype space. A member of the population with a genotype $x \in \{0, 1\}^N$ is replaced in one time step by a fitter one-mutation neighbour with genotype y with probability proportional to increase in fitness (with regard to f) and normalised by the other possible evolutionary steps. That is, if a member has genotype x , then the probability that it has a different genotype, y , in the next time step

$$\Pr(x \rightarrow y) = \begin{cases} \frac{f(y)-f(x)}{\sum_{Ham(x,z)=1} \max\{f(z)-f(x), 0\}} & f(y) > f(x) \\ 0 & \text{otherwise} \end{cases}. \quad (2)$$

If we further define, for each genotype x , a probability

$$\Pr(x \rightarrow x) = \begin{cases} 1 & \text{if } x \text{ has no fitter one-step neighbours} \\ 0 & \text{otherwise} \end{cases}, \quad (3)$$

our model becomes a time-homogeneous absorbing Markov chain with a finite state space and transition matrix $P = [p_{ij}]$ where $p_{ij} = \Pr(i \rightarrow j)$ for each $i, j \in \{0, 1\}^N$.

Using this Markov chain we can explore how a population of a given organism comprising a variety of genotypes evolves over time. To do this we make the assumption that the population size is large and remains constant so that only beneficial mutations will fix in the population. We define a collection of population row vectors $\mu^{(t)}$ for each $t \in \mathbb{N}$, where $\mu^{(t)}$ is a vector of length 2^N in which $\mu_k^{(t)}$ is the proportion of the population in the k th genotype at time t . Given such a population at time t the one-step update can be computed by

$$\mu^{(t+1)} = \mu^{(t)} P. \quad (4)$$

As a result of the associativity of matrix multiplication we can compute the distribution of a population at time t given an initial population $\mu^{(0)}$ by

$$\mu^{(t)} = \mu^{(0)} P^t. \quad (5)$$

As we have now encoded the same evolutionary process in two different formalisms, one stochas-




Figure 3: A comparison of the two models. For a given randomly generated fitness landscape over a genotype space with $N = 5$ we plot the distribution of the population at the 7 possible peaks when each algorithm terminates. The dark blue bars represent the average of 1000 runs of the explicit algorithm with a starting population of size 30. The error bars give the standard error in the predictions of this model. The light blue bars give the population distribution predicted by the Markov chain model

tic and one deterministic, we compare the results of each with identical initial conditions. Figure 3 shows an example output for each model given a single, randomly generated landscape over a genotype space with $N = 5$. The starting populations, upon which the ending populations are highly dependent, are identical. In this case, the stochastic process was run until termination in all cases, and the deterministic model was calculated to steady state. We find that the results of the Markov Chain model fall, as expected, within the standard error of the stochastic model.

Evolution on Large Time Scales

To explore long term evolutionary properties of populations using an explicit evolutionary algorithm we would have to iterate the population for a large number of generations. For large populations, or for ones with realistic genotype lengths, this can be prohibitively inefficient. In our Markov model of evolution we can simulate each time step much more efficiently and so can explore the long term behaviour much more effectively. While this increase in computational efficiency is a benefit of the

model, a more significant improvement is revealed by its algebraic structure. As our evolutionary process is now encoded by a Markov Chain we can explore properties of the process analytically by examining the transition matrix P with no need for simulations at all. The following lemma explores what happens in the evolutionary process over large time scales.

Lemma:

Let P be a transition matrix as described above. Then there exists some k such that $P^k P = P^k$

Proof:

We first note that, without loss of generality, we may (by rearranging the population vector ordering) assume that our transition matrix P has the block matrix form:

$$P = \begin{bmatrix} Q & R \\ 0 & I \end{bmatrix}, \quad (6)$$

where Q is a transition matrix encoding the probabilities of moving between non-absorbing (i.e. non fitness optima) genotypes, R encodes the probabilities of moving from non-absorbing states into absorbing states, and I is the identity matrix encoding that once in an absorbing state the walk will remain there. Now taking powers of P gives

$$P^k = \begin{bmatrix} Q^k & (I + Q + Q^2 + \dots + Q^{k-1})R \\ 0 & I \end{bmatrix}. \quad (7)$$

Note that if $Q^k = 0$ for some k then $P^k P = P^k$ as

$$P^k P = \begin{bmatrix} 0 & (I + Q + Q^2 + \dots + Q^{k-1})R \\ 0 & I \end{bmatrix} \begin{bmatrix} Q & R \\ 0 & I \end{bmatrix} = \begin{bmatrix} 0 & (I + Q + Q^2 + \dots + Q^{k-1})R \\ 0 & I \end{bmatrix} = P^k \quad (8)$$

and so the lemma follows if $Q^{2^N+1} = 0$. Assume to the contrary that $Q^{2^N+1} \neq 0$. Writing $Q^{2^N+1} = [q_{i,j}^*]$ we have for some n, m that $q_{n,m}^* \neq 0$. By the definition of the power of a stochastic matrix the probability of transitioning from the n th to m th non-absorbing states in $2^N + 1$ steps is non-zero. However, we have by construction that any evolutionary paths amongst those genotypes which are non-absorbing are increasing in fitness at every step and are necessarily acyclic. It follows then, that as $q_{n,m}^* \neq 0$ there exists an acyclic path of length $2^N + 1$ in a space of fewer than 2^N genotypes. This yields a contradiction and the lemma follows. ■

As a consequence we know that the matrix

$$P^* = \lim_{k \rightarrow \infty} P^k \quad (9)$$

exists and in fact this limit is found after only finitely many iterations. It follows that a given initial population $\mu^{(0)}$ will converge to a stationary distribution μ^* after a finite number of steps in our evolutionary model. Furthermore if P^* is known then we can determine this stationary distribution by the calculation

$$\mu^* = \mu^{(0)} P^*. \quad (10)$$

Therefore, the limit matrix P^* need only be calculated once and can be used to explore the behaviour of any number of different starting configurations. This offers an improvement over an explicit evolutionary model as we can compute the limit matrix P^* by repeatedly squaring the transition matrix P . The previous lemma shows we need to square P at most N times to find P^* . We know that the product of two $m \times m$ matrices can be computed in time $\mathcal{O}(m^{2.807})$ by Strassen's algorithm (Strassen [1969]). It follows that this simulation has worst case time complexity $\mathcal{O}(N 2^{2.807N})$ in computing P^* , although in many cases it is considerably faster. This is slower than a single run of the explicit simulation, but considerably faster than using the explicit simulation to determine the expected population distribution.

It could be argued that this construction introduces determinism into the inherently stochastic process of evolution but this is not so. The population vector μ^* represents an expected population distribution after a large number of iterations and gives the same information as running an explicit simulation many times and averaging the results (cf. Fig. 3).

Discussion

We have taken an explicit evolutionary algorithm and encoded it as an equivalent Markov process which reduces each update step of the algorithm to a single matrix multiplication. Further we have seen that we can reduce the problem of determining the evolutionary trajectory of a given starting population to a single matrix multiplication with the matrix P^* which can be efficiently computed.

There is a rich theory of how the form of a matrix determines its properties under multiplication and exponentiation and the previous examples show that these properties can help us gain useful insight into the evolutionary process by using familiar analytical tools from linear algebra. In particular it is of interest to ask how different assumptions about our evolutionary process change the matrix P and hence what predictions we can make about a population. An example might be

to ask how does allowing mutations through neutral spaces (those for which the fitness function f remains constant) affect the evolutionary dynamics? It has been established that neutral spaces might have a significant impact on how populations evolve (Schaper et al. [2012]). Certainly we lose the fact that the matrix P^* will be found after a finite number of iterations of the process - there exist infinite walks in our Markov Chain provided a neutral space exists.

Further, as Tan and colleagues showed (Tan et al. [2011]), certain paths in phenotype space can be obviated by changing landscapes. This algebraic construction could be used to analytically design landscapes to effectively steer evolution through these high-dimensional spaces, offering the possibility of new uses for old drugs to help avoid the evolution of resistance in pathologic states such as cancer or infection.

Acknowledgment and Supplementary Information

The authors would like to thank Arne Traulsen at the Max Planck Institute for helpful discussions. We would like to offer the code underlying this model to any interested parties openly. If interested, please email DN.

Bibliography

- P.A.P. Moran et al. The statistical processes of evolutionary theory. *The statistical processes of evolutionary theory.*, 1962.
- D.M. Weinreich, R.A. Watson, and L. Chao. Perspective: sign epistasis and genetic constraint on evolutionary trajectories. *Evolution*, 59(6):1165–1174, 2005.
- Daniel M Weinreich, Nigel F Delaney, Mark A Depristo, and Daniel L Hartl. Darwinian evolution can follow only very few mutational paths to fitter proteins. *Science*, 312(5770):111–4, Apr 2006. doi: 10.1126/science.1123539.
- L. Tan, S. Serene, H.X. Chao, and J. Gore. Hidden randomness between fitness landscapes limits reverse evolution. *Physical Review Letters*, 106(19):198102, 2011.
- F.J. Poelwijk, S. Tănase-Nicola, D.J. Kiviet, and S.J. Tans. Reciprocal sign epistasis is a necessary condition for multi-peaked fitness landscapes. *Journal of Theoretical Biology*, 272(1):141–144, 2011.
- Richard W. Hamming. Error detecting and error correcting codes. *Bell System Technical Journal*, 29(2):147–160, 1950.
- R. A. Fisher. *The genetical theory of natural selection*. Clarendon Press; Oxford University Press., 1930.
- Lee Altenberg. Nk fitness landscapes. *Handbook of Evolutionary Computation*, 1997.
- V. Strassen. Gaussian elimination is not optimal. *Numerische Mathematik*, 13(4):354–356, 1969.
- Steffen Schaper, Iain G Johnston, and Ard A Louis. Epistasis can lead to fragmented neutral spaces and contingency in evolution. *Proc Biol Sci*, 279(1734):1777–83, May 2012. doi: 10.1098/rspb.2011.2183.